arXiv:1611.05216v1 [cs.CV] 16 Nov 2016

shuttleNet: A biologically-inspired RNN with loop connection and parameter sharing

Yemin Shi[1,2], Yonghong Tian[1,2]*, Yaowei Wang[3]*, Tiejun Huang[1,2]
[1] National Engineering Laboratory for Video Technology, School of EE&CS,
Peking University, Beijing, China
[2] Cooperative Medianet Innovation Center, China
[3] School of Information and Electronics,
Beijing Institute of Technology, Beijing, China

Abstract

Despite a lot of research efforts devoted in recent years, how to efficiently learn long-term dependencies from sequences still remains a pretty challenging task. As one of the key models for sequence learning, recurrent neural network (RNN) and its variants such as long short term memory (LSTM) and gated recurrent unit (GRU) are still not powerful enough in practice. One possible reason is that they have only feedforward connections, which is different from biological neural network that is typically composed of both feedforward and feedback connections. To address the problem, this paper proposes a biologically-inspired RNN structure, called shuttleNet, by introducing loop connections in the network and utilizing parameter sharing to prevent overfitting. Unlike the traditional RNNs, the cells of shuttleNet are loop connected to mimic the brain's feedforward and feedback connections. The structure is then stretched in the depth dimension to generate a deeper model with multiple information flow paths, while the parameters are shared so as to prevent shuttleNet from being overfitting. The attention mechanism is then applied to select the best information path. The extensive experiments are conducted on two datasets for action recognition: UCF101 and HMDB51. We find that our model can outperform LSTMs and GRUs remarkably. Even only replacing the LSTMs with our shuttleNet in a CNN-RNN network, we can still achieve the state-of-the-art performance on both datasets.

1. Introduction

Deep neural networks (DNNs) have achieved a great success in recent years. DNNs are able to learn complex features and patterns from raw data. A typical DNN has multiple nonlinear layers which are connected with feedforward connections. However, there are a lot of limitations in feedforward connections. For example, the feedforward connection brings a problem that in order to perform better, DNN has to contain a lot of layers, which will lead to massive parameters and make it easy for DNN to be over-fitting.

Overall speaking, DNNs are a kind of brain-inspired deep network architecture. The visual areas of the brain are interconnected in a complex pattern of feedforward, lateral, and feedback pathways [10, 24]. Feedback connections are ubiquitous throughout the cortex, and subcortical regions in ascending hierarchical pathways also receive a large amount of feedback from cortical areas [9, 11, 13, 29]. This motivated us that DNNs may benefit a lot from imitating the brain connection patterns.

Following the similar idea, recurrent neural networks (RNNs) introduce temporal connections in the network to condition their present state on their entire history of inputs. Because of RNNs' recurrent lateral connection mechanism, they are able to capture long-term dependencies in sequential data over an extended period of time. Moreover, RNNs have been theoretically proved to be a Turing-complete machine, which means they can be used to approximate any functions [32]. As one variant of RNNs, long short term memory (LSTM) [15] was proposed to solve the gradient vanishing and exploding problems. When unfolded in time, LSTMs are equivalent to very deep neural networks that share model parameters and receive the input at each time step. The parameter sharing guarantees that there will not be too many parameters and consequently the network is trainable. Many works [7, 30, 26, 39, 40, 52, 53] have proved the effectiveness of LSTM on the action recognition task from video sequences.

In this paper, we propose a more powerful RNN structure, shuttleNet, by introducing loop connections in the network and utilizing parameter sharing to prevent overfitting.

*Corresponding author: Yonghong Tian (email: yhtian@pku.edu.cn) and Yaowei Wang (yaoweiwang@bit.edu.cn).

In our shuttleNet, multiple RNN cells are organized into a ring, and then the loop connections are generated by rotating the ring every several steps and connecting the cells at the same position after rotating. We call this process as *depth stretching*, which will result in a deeper model with multiple information flow paths. In order to restrict the number of parameters, we share all parameters between different rotations while keeping an individual hidden state at each rotation. We also employ the attention mechanism [2] to help select the best information flow path. As a result, the shuttleNet is expected to achieve better performance, without the risk of over-fitting.

We evaluate the proposed shuttleNet for the action recognition task on two popular datasets, namely UCF101 and HMDB51. Experimental results show that the shuttleNet outperforms the LSTMs and GRUs remarkably, and yields the state-of-the-art performance on both datasets.

The rest of the paper is organized as follows: In section 2, we review the related work. The proposed shuttleNet is presented in section 3. Experimental results are discussed in section 4. Finally, section 5 concludes this paper.

2. Related work

Basically, action recognition aims at categorizing the actions or behaviors of one or more persons in a video sequence. Modeling and utilization of the long-term dependence in videos has been proven effective to improve the performance significantly by many works. Two-stream ConvNets [33] is widely recognized as the first successful deep learning framework for action recognition. It extracts the spatial and temporal characteristics in one framework, and trains the standalone CNNs for two streams separately. Wang *et al.* [48] also successfully trained very deep two-stream ConvNets on the UCF101 dataset. Similarly, trajectory-pooled deep-convolutional descriptor (TDD) was proposed by Wang *et al.* [46], which shares the merits of both hand-crafted features such as dense trajectories [41, 43] and deeply-learnt features. However, two-stream ConvNets did not capture and utlize the long-term dependence in the network.

Hierarchical recurrent neural network [8] is one of the earliest works which attempted to improve the efficiency of capturing long term dependency. Long short term memory (LSTM) [15], the most successful approach to deal with vanishing gradients till to now, was proposed by Hochreiter and Schmidhuber. Basically, LSTM relies on a fantastic structure made of gates to control the flow of information to the hidden neurons. Peephole LSTM [12] adds peepholes to some gates so as to allow them look at the cell state. Gated recurrent unit (GRU), introduced by Cho *et al* [4], is a slightly more dramatic variation on the LSTM, which combines the forget and input gates into a single update gate and merges the cell state and hidden state. Mikolov *et al.*

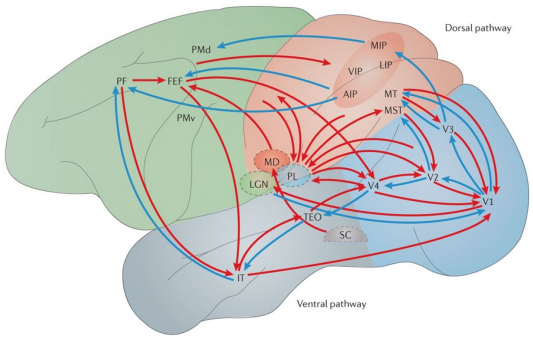

Figure 1. Illustration of the visual cortical pathways [17]. Feedforward connections are represented by blue arrows and feedback connections are represented by red arrows.

[25] proposed to add a hidden layer to RNNs and make the weight matrix close to identity. Clockwork RNN, proposed by Koutnik *et al.* [19], splits each hidden layer into several modules running at different clocks and outputting at the predefined clock rate. Gated feedback recurrent neural networks [5] attempt to allow the model to adaptively adjust the connection between consecutive hidden layers by using the gated feedback connection between layers of stacked RNNs.

It should be noted that LSTMs were introduced to model long-term actions for action recognition recently. Yue-Hei *et al.* [53] and Donahue *et al.* [7] proposed their own recurrent networks respectively by connecting LSTMs to CNNs. Wu *et al.* [51] achieved the state-of-the-art performance by connecting CNNs and LSTMs under the hybrid deep learning framework. Shi *et al.* [31, 30] also introduced their DTD and sDTD to model the dependence on the temporal domain. Nevertheless, LSTMs and GRUs are still not powerful enough for action recogntion in practice. One possible reason is that they have only feedforward connections, which is different from biological neural network that is typically composed of both feedforward and feedback connections. Therefore, this paper proposes a biologically-inspired RNN structure, called shuttleNet, by introducing loop connections in the network. We will show the shuttleNet outperforms the LSTMs and GRUs remarkably for action recognition.

3. shuttleNet

In this section, we will first explain our basic motivation from the inconsistency between the brain connection patterns and DNNs, then present the overall framework of the proposed shuttleNet. After that, we will describe the details about its key components one by one.

3.1. Motivation

Hierarchical processing is the key to understanding vision system. It consists of hierarchically organized dis-

Table 1. Some popular deep networks and their depths.

LeNet-5	VGG-16	GoogLeNet	ResNet-152
5	16	22	152

Please refer to [22, 34, 37, 14].

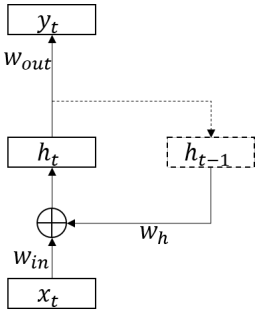

Figure 2. Illustration of the recurrent neural network (RNN). The dash line means h_{t-1} comes from the previous time step.

tinct anatomical areas functionally specialized for processing different aspects of a visual object [10]. These visual areas are interconnected through ascending feedforward connections, descending feedback connections and connections from neural structures at the same hierarchical level [21]. The lateral geniculate nucleus (LGN) is a relay center in the thalamus for the visual pathway. There is evidence that only 10% of inputs to LGN come from the retina and 90% are feedback modulatory inputs from cortex and the brainstem. As shown in Figure 1, the feedback connections play an important role in visual cortical pathways. When ignoring the function of feedforward and feedback connections, we can simplify the visual pathways as multiple loop flow where the vision information is transformed along.

On the contrary, DNN researchers concentrate on designing deeper network with only feedforward connections, as shown in Table 1. As the model becoming more and more deeper, it will consume more time and computing resource to train. However, LGN and V1 have only 6 layers but still are powerful enough for understanding vision information. A reasonable way to re-design the networks is to imitate the feedback connections of visual cortical pathways.

3.2. The overall framework of shuttleNet

To begin the discussion of shuttleNet, we first simply review the formulation of RNN. Essentially, an RNN is a network which learns how to predict based on the previous state and current input. At each time step t, for an input x_t, an RNN updates the hidden state recursively as follows (as shown in Figure 2):

$$h_t = f(w_i x_t + w_h h_{t-1}) \tag{1}$$

where $f(\cdot)$ is a nonlinear activation function, mostly the hyperbolic tangent nonlinearity defined as $\phi(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$, and w_i and w_h are the weight matrices for the input and hidden states.

Given the hidden state h_t, RNN generates the predictions according to the following equation:

$$y_t = g(w_o h_t) \tag{2}$$

where $g(\cdot)$ denotes the softmax function and w_o is the weight matrix for the output. Typically, RNN is trained with the back-propagation through time (BPTT) algorithm [50].

Basically, shuttleNet is a biologically-inspired RNN structure by introducing loop connections in the network and utilizing parameter sharing to prevent overfitting. As shown in Figure 3, the input is replicated n times so that each RNN cell has its own input. At every round, all RNN cells take their corresponding inputs and generate outputs while maintaining a hidden state. After each round, the RNN cells will rotate k steps. In the depth dimension, RNN cells work at all rounds to share the weights. Finally, based on RNN cells' last outputs and the input, the attention mechanism is applied to select the best path as the final output.

A noticeable advantage of shutlleNet is that, even though having a complicated connection, the number of parameters does not increase. This can effectively prevent shuttleNet from over-fitting. Moreover, because every RNN cell works at multiple network depths, the gradients from different depths will force the RNN cells to simultaneously consider how to benefit all depths.

It should be noted that when n and d equal 1, shuttleNet is completely equivalent to a RNN cell. When n is greater than 1 and d equals 1, shuttleNet works like the model ensemble while having the ability to automatically select the best model. When n equals 1 and d is geater than 1, shuttleNet is a weights-shared stacked RNN. If the weights are not shared in the depth dimension, shuttleNet is similar to the stacked RNN.

3.3. Loop connection

As discussed in section 3.1, the visual pathways have more feedback connections than feedforward connections. Our first intuition is to design a network with feedforward and feedback connections working like the visual pathways. However, to make the model computable and stable, we choose to implement the connections as a simplified version, loop connections. We will then describe it with mathematical linguistic forms.

For most known network structure, layers are stacked one by one and the last layer produces the final output. The layers of a network can be considered as some nonlinear functions. For input i, a two-layer network generate the output as follows:

$$o = f(g(i)) \tag{3}$$

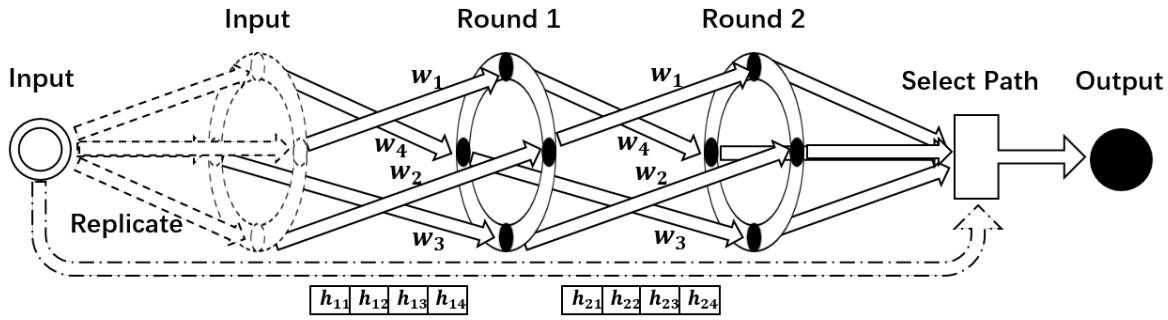

Figure 3. Illustration of a 4-cell-2-round shuttleNet. The input is replicated so that each unit has its own input. All units are formed as a circle and work at every round simultaneously, while all rounds share the same weights. The units have individual hidden states at each round. The information flow is propagated in a wheel way (rotate k steps after one round). The attention mechanism is utilized to select the best flow path.

However, there is no rule claiming that we should apply $f(\cdot)$ first or $g(\cdot)$ first. Thus we may get a better performance with:

$$o = g(f(i)) \tag{4}$$

In the general case, we can never determine the best order for each sample. A brute-force solution is to iterate through all the possibilities which is very computational expensive. Our solution is organizing all layers into a ring. The input is first replicated and fed into all layers (as shown in Figure 3, the dash line ring). Then we decide the top layer for each layer by rotating the ring k steps, where k is a fixed integer. The output of one layer is fed into the layer at the same position after rotating.

Formally, for n layers $(f_0, f_1, ..., f_{n-1})$ and a given input i, we compute n outputs as follows:

$$o_j = f_{(j+d*k)\%n}(...(f_{(j+k)\%n}(f_j(i)))) \tag{5}$$

where $\%$ is the modulus operator, d is the user defined maximum nesting depth and o_j is the output of layer j. Therefore, there are totally n outputs which are then used to generate the final output.

3.4. Depth stretching

For fully-connected layers, the loop connection should be able to work correctly. However, for a recurrent neural network, the hidden state is used to remember previous information. The u-th RNN cell r_u at v-th round t-th time step in the loop connection works as follows:

$$o^t_{u,v}, h^t_u = r_{u,v}(o^t_{u-k,v-1}, h^t_u) \tag{6}$$
$$o^t_{u,v+1}, h^t_u = r_{u,v+1}(o^t_{u-k,v}, h^t_u) \tag{7}$$

where $o^t_{u,v}$ is the output of u-th RNN cell at v-th round and h^t_u is the hidden state of u-th RNN cell. As we can see, the hidden state is overwritten after each round. At the meantime, the $o^t_{u,v+1}$ will rely on h_u which is computed at a different network depth. And $o^{t+1}_{u,1}$ will also rely on h_u which is computed together with $o^t_{u,d}$. The chaotic dependencies will end up with unstable training process.

We propose to keep an individual hidden state for each RNN cell at every round while sharing all weights. In this way, we can stretch the loop connection in the depth dimension (as shown in Figure 3). The u-th RNN $r_{u,v}$ at v-th round in the depth stretched structure works as follows:

$$o^t_{u,v}, h^t_{u,v} = r_{u,v}(o^t_{u-k,v-1}, h^t_{u,v-1}) \tag{8}$$
$$o^t_{u,v+1}, h^t_{u,v+1} = r_{u,v+1}(o^t_{u-k,v}, h^t_{u,v}) \tag{9}$$

where $h^t_{u,v}$ is the hidden state of u-th RNN cell at v-th round t-th time step.

3.5. Path selecting

According to section 3.3 and 3.4, given input i^t and n RNN cells at the t-th time step, we get n outputs $(o^t_{1,d}, o^t_{2,d}, ..., o^t_{n,d})$. We will have to decide which RNN cell path will produce the best prediction. In this paper, the attention mechanism [2] is employed to help select the proper output.

To compute the attention vector at each output time t over the n outputs, we define:

$$e^t_j = \nu^T tanh(w'_i i^t + w'_o o^t_{j,d}) \tag{10}$$
$$\alpha^t_j = \frac{exp(e^t_j)}{\sum_{l=1}^{n} exp(e^t_l)} \tag{11}$$
$$O^t = \sum_{j=1}^{n} \alpha^t_j o_j \tag{12}$$

where vector ν and the weight matrixes w'_i, w'_o are the parameters of the model and the O^t is the output of shuttleNet. The vector e^t assigns a weight for each output $o^t_{j,d}$, which means how much attention should be put on $o^t_{j,d}$. These attention weights are normalized by softmax to create the attention mask α^t over the outputs.

4. Experiments

In this section, we will first introduce the detail of datasets and their corresponding evaluation schemes. Then, we describe the implementation details of our model. To find out the effect of each parameter, we explore experiments with multiple parameter settings and prove the effectiveness of loop connection and parameter sharing. We finally report the experimental results and compare shuttleNet with the state-of-the-art methods to demonstrate the superior performance of shuttleNet.

4.1. Datasets

To verify the effectiveness of shuttleNet, we conduct experiments on two public datasets: HMDB51 [20] and UCF101 [35].

The HMDB51 dataset is a large collection of realistic videos from various sources, including movies and web videos. It is composed of 6,766 video clips from 51 action categories, with each category containing at least 100 clips. The action categories include simple facial actions, general body movements and human interactions. Our experiments follow the original evaluation scheme, and average accuracy over the three train-test splits is reported.

The UCF101 dataset is one of the most popular action recognition benchmarks. It contains 13,320 video clips (27 hours in total) from 101 action classes and there are at least 100 video clips for each class. The 101 classes are divided into five groups: Body-Motion, Human-Human Interactions, Human-Object Interactions, Playing Musical Instruments and Sports. Following [16], we conduct evaluations using 3 train/test splits, which is currently the most popular setting in using this dataset. Results are measured by classification accuracy on each split and we report the mean accuracy over the three splits.

Compared with the very large datasets used for image classification, the datasets for action recognition is relatively smaller. Therefore, we pre-train our model on the ImageNet dataset [6]. As UCF101 is larger than HMDB51, we also use UCF101 to train our model initially, and then transfer the learnt model to HMDB51.

4.2. Implementation details

Our shuttleNet is tested on action recognition task. As shown in Figure 4, we use a CNN-RNN network structure as in [7, 30] and use two-stream framework [33] to get our final prediction. The video frames are first fed into CNN to learn representations. The CNN representations are used by RNN to learn temporal features. Then the output of RNN is used to predict action label. In this paper, the CNN part is implemented with GoogLeNet [37] without the auxiliary losses. The RNN part is tested with several options: 2-layer LSTM[15], 2-layer GRU[4] and our shuttleNet.

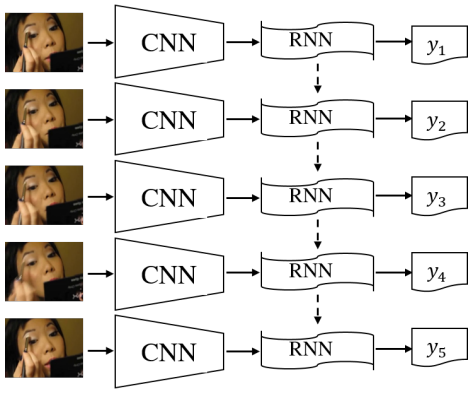

Figure 4. Illustration of our action recognition network structure. The video frames are fed into CNN to learn representations. The CNN representations are then fed into RNN to learn temporal features. Finally, the output of RNN is used to predict action label.

We implement shuttleNet with TensorFlow [1] and use GRU as the RNN cell just because of less parameter number and good performance compared with other RNNs.

To extract optical flow, we choose the TVL1 optical flow algorithm [54] and use the OpenCV GPU implementation. We discretize the optical flow fields into interval of $[0, 255]$ by a linear transformation and saved them as images. To further improve the performance, we also employ the warped TVL1 optical flow field [49]. Unlike the two-stream ConvNets in [33], whose temporal stream input is volumes of stacking optical flow fields ($224 \times 224 \times 2F$, where F is the number of stacking flows and is set to 10), our temporal stream input is single optical flow. An optical flow field is computed from two consecutive frames and composed of vertical and horizontal flows. To make use of the optical flow, the flow-x, flow-y and their quadratic mean is used to form a three channel image.

After initializing with the pre-trained ImageNet model, the network is trained with mini-batch stochastic gradient descent with momentum (set to 0.9). We read 16 frames/optical flows with a stride of 5 from each video as one sample for the RNN. We resize all input images to 256×256, and then use the fixed-crop strategy [48] to crop a 224×224 region from images or their horizontal flip. Because the 16 consecutive samples are needed in the RNN, we also force images from the same video to crop the same region. In the test phase, we sample 4 corners and the center from each image and its horizontal flip, and 25 samples are extracted from each video. The batch size is 16 (videos) and the learning rate starts from 0.01. For spatial stream, the learning rate is divided by 10 at iteration $10K$, $20K$ and $30K$, and training is stopped at $40K$ iteration. For temporal stream, the learning rate is divided by 10 at $20K$, $30K$ and $40K$, and training is stopped at $45K$ iteration.

In the remainder of the paper, we use spatial stream and temporal stream to indicate the streams in two-stream

Table 2. Accuracy of shuttleNet with different parameter setting on UCF101 split 1.

No. RNN cells	No. rounds	Spatial	Temporal
2	1	82.37%	82.76%
1	2	82.9%	82.92%
2	2	83.03%	82.18%
2	3	80.89%	82%
3	2	83.19%	82.66%
3	3	82.13%	80.84%

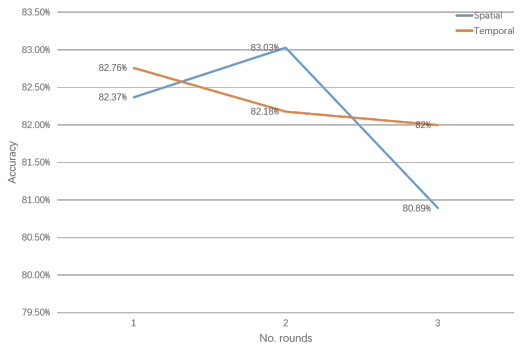

Figure 5. The performance under different round and 2 RNN cells.

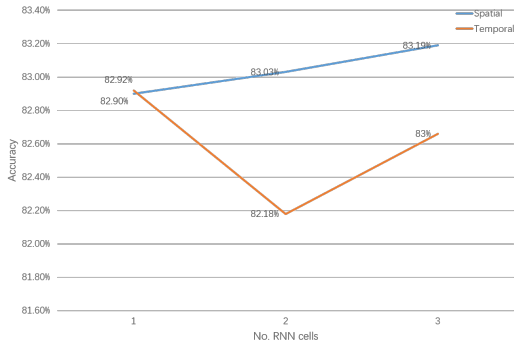

Figure 6. The performance under different number of RNN cells and 2 rounds.

framework and use warped temporal stream to indicate the stream whose input is warped TVL1 optical flow.

4.3. Evaluation

In this section, we will first test shuttleNet with different setting by adjusting the number of RNN cells and the number of rounds in section 4.3.1. We also compare our model with several existing RNN models in section 4.3.2. We will prove that our shuttleNet can achieve much better performance with similar number of parameters. Finally, according to previous experiments, we test our model on more datasets with the best setting in section 4.3.3.

Table 3. Comparison of shuttleNet to several popular recurrent neural networks on UCF101 split 1.

Model	#	Spatial	Temporal	shuttleNet-2†
GoogLeNet‡	-	77.74%	79.75%	85.91%
LSTM [15]	2	78.8%	80.15%	88.58%
GRU [4]	2	80.49%	82.77%	90.27%
GRU	3	82.42%	81.21%	90.64%
shuttleNet	1	82.9%	82.92%	91.65%
shuttleNet	2	83.03%	82.18%	91.94%

† Merge the spatial stream and the temporal stream [33].
‡ No RNN in GoogLeNet [37].
The second column is the number of the RNN cells.

4.3.1 Exploration experiments

There are two pre-defined parameters: the number of RNN cells n and the number of rounds d. We will first explore the effect of each parameter by running experiments under multiple parameter setting combinations.

The performances of shuttleNet under different parameters are shown in Table 2. We plot the accuracy of shuttleNet with different round and 2 RNN cells in Figure 5. We can easily find that as the number of rounds growing, the performance of temporal stream first increases then decreases and the performance of spatial stream keeps decrease slightly. As shown in Figure 6, as the number of RNN cells growing, the temporal stream performance first decreases then increases and the spatial stream keeps increase.

Actually, the above conclusions are very easy to understand. When increasing the number of rounds, each cell will have to take care of more situations. The more rounds used, the harder for RNN cells to learn effective feature. However, with 2 rounds, it works like regularization approach preventing overfitting hence improving the performance. When increasing the number of RNN cells, all cells work like model ensemble. The more cells used, the better the ensemble will be.

Based on the above discussion, we will use 2 rounds in the following sections. Although our shuttleNet can be easily extended as a multi-layer model, we use only one layer shuttleNet with 2 RNN cells to compare with other models, for example, 2-layer LSTM or GRU, to emphasize the good performance and ensure fairness.

4.3.2 Comparison with baselines

As shown in Table 3, we compare our shuttleNet with several popular recurrent neural network model. Our model achieves a good performance of 91.94% on UCF101 split 1 while none of other models close to 91%. On the spatial stream, shuttleNet outperforms other RNN models remarkably. Even the three layer GRU network can not beat our

Table 4. Number of parameters in 2-layer RNN models with an output size of 1024.

Model	LSTM	GRU	shuttleNet
#	16,777,216	8,388,608	10,485,760

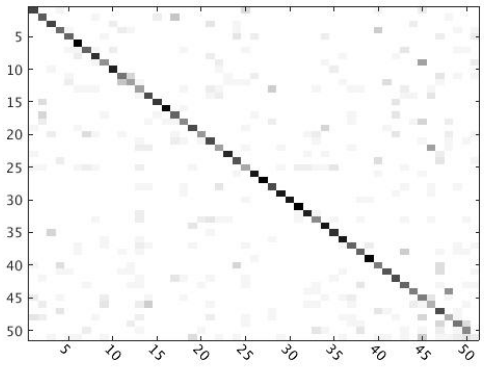

Figure 7. Performance confusion matrix on the HMDB51 split 1.

2-cell shuttleNet.

To prove the effectiveness of the loop connection, we conduct an experiment with a 2-cell-1-round shuttleNet. The accuracy of spatial stream is 82.37% and temporal accuracy is 82.76%. The shuttleNet-2 achieves an accuracy of 91.7%. This means that loop connection is very effective and can improve the performance significantly.

It's amazing that our 1-layer-2-round shuttleNet also achieve an accuracy of 91.65%, which is still much better than all other RNN models. The great performance of 1-layer shuttleNet proves the effectiveness of our parameter sharing mechanism.

The parameter sharing mechanism also makes sure that our model won't have too many parameters. As shown in Table 4, our model has slightly more parameters than GRU while having much less parameters than LSTM. The addition parameters are used to compute the attention value for each path. This means that no matter how many RNN cells we have, there won't be any more parameters needed.

Parallel computing is another advantage of our shuttleNet. Because that each cell in shuttleNet at every round works independently, it's easy to accelerate the computing by parallel computing all cells.

4.3.3 Evalution of shuttleNet

From section 4.3.2, we prove the effectiveness of shuttleNet on UCF101 split 1. We report the performance of shuttleNet on more splits and dataset in Table 5.

For the HMDB51 dataset, shuttleNet perform good on the first split while relative worse on the other two splits. And the overall performance is not as good as in UCF101.

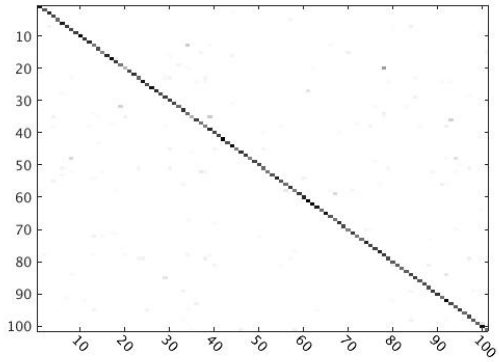

Figure 8. Performance confusion matrix on the UCF101 split 1.

We believe this is because of the small training set size (3570 videos). But shuttleNet is still outstanding on HMDB51 when comparing to other methods.

Our shuttleNet achieve the best performance on UCF101 split 3. This should be related to its biggest train set. The mean accuracy of shuttleNet-2 over the three splits is 92.29% and our shuttleNet-3's mean accuracy is 92.55%. We achieve a surprisingly high accuracy of 87.74% on UCF101 split 3 with temporal stream, but get a relative lower accuracy with spatial stream. We think that's because of the bad split and bad frame quality or more highly motion related actions.

The confusion matrixes for shutttleNet on HMDB51 and UCF101 datasets are shown in Figure 7 and 8. The confusion matrix on the UCF101 dataset is well diagonalized. However, the confusion matrix on the HMDB51 dataset shows that some categories are easily mis-classified, despite shuttleNet still performs well on most categories.

4.3.4 Comparison with the state-of-the-art methods

After exploring of the practices and understanding the effect of shuttleNet, we are ready to build up our final action recognition method. Specifically, we assemble three input modalities and all the techniques described as our final recognition approach, and test it on two challenging datasets: HMDB51 and UCF101. The results are summarized in Table 6, where we compare our method with both traditional approaches such as iDT and deep learning method such as hybrid deep framework [51], very deep two-stream [48].

Even though we use a very simple framework and only replace LSTM with our shuttleNet, our model still outperforms these methods on UCF101 and surpasses most methods on HMDB51. The superior performance of our methods demonstrates the effectiveness of shuttleNet.

Table 5. Accuracy of shuttleNet on two datasets.

Dataset	split	Spatial	Temporal	Warp Temporal	shuttleNet-2	shuttleNet-3‡
HMDB51	1	50.07%	53.53%	54.77%	65.36%	66.28%
	2	46.21%	55.1%	55.82%	61.83%	62.68%
	3	45.69%	56.6%	56.8%	63.66%	64.12%
	Mean	47.33%	55.08%	55.8%	63.62%	**64.36%**
UCF101	1	83.03%	82.18%	81.26%	91.94%	92.15%
	2	80.21%	84.73%	82.14%	92.21%	92.66%
	3	80.95%	87.74%	84.39%	92.72%	92.83%
	Mean	81.4%	84.88%	82.6%	92.29%	**92.55%**

‡ Merge the spatial, temporal and warp temporal streams [49].

Table 6. Comparison of sDTD to the state-of-the-art methods.

HMDB51		UCF101	
STIP+BoF [20]	23.0%	STIP+BoF [20]	43.9%
Motionlets [45]	42.1%	Deep Net [18]	63.3%
DT+BoF [42]	46.6%	DT+VLAD [3]	79.9%
DT+MVSV [3]	55.9%	DT+MVSV [3]	83.5%
iDT+FV [43]	57.2%	iDT+FV [44]	85.9%
iDT+HSV [27]	61.1%	iDT+HSV [27]	88.0%
PMF+AdaBoost+Correlogram+SVM [23]	36.5%	Hybrid deep framework [51]	91.3%
Two-stream model [48]	59.4%	Two-stream model [48]	88.0%
F_{ST}CN [36]	59.1%	F_{ST}CN [36]	88.1%
TDD+FV [46]	63.2%	TDD+FV [46]	90.3%
TDD+iDT+FV [46]	**65.9%**	TDD+iDT+FV [46]	91.5%
Visual Attention [28]	41.3%	Very deep two-stream [48]	91.4%
MoFAP [47]	61.7%	MoFAP [47]	88.3%
LTC [38]	64.8%	LTC [38]	91.7%
-	-	LSTM with 30 Frame Unroll [53]	88.6%
shuttleNet-2	63.62%	shuttleNet-2	92.29%
shuttleNet-3	**64.36%**	shuttleNet-3	**92.55%**

5. Conclusion

In this paper, we present a biologically-inspired RNN, shuttleNet, which has a special loop connection and shares parameters in the depth dimension. As demonstrated on two challenging datasets, shuttleNet can outperform other RNN models like LSTM and GRU remarkably while having roughly equal number of parameters. This is largely ascribed to the loop connections as well as parameter sharing mechanism. The former breaks the limitation of feedforward connections, while the latter prevent shuttleNet from being over-fitting.

There are still a lot of research directions that will be addressed in the future, for example, to check the performance of the stacked shuttleNet and whether CNN will benefit from loop connection.

References

[1] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et al. Tensorflow: Large-scale machine learning on heterogeneous distributed systems. *arXiv preprint arXiv:1603.04467*, 2016. 5

[2] D. Bahdanau, K. Cho, and Y. Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014. 2, 4

[3] Z. Cai, L. Wang, X. Peng, and Y. Qiao. Multi-view super vector for action recognition. In *Proceedings of IEEE conference on Computer Vision and Pattern Recognition*, pages 596–603, 2014. 8

[4] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*, 2014. 2, 5, 6

[5] J. Chung, C. Gülçehre, K. Cho, and Y. Bengio. Gated feedback recurrent neural networks. *CoRR, abs/1502.02367*, 2015. 2

[6] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database.

In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 5

[7] J. Donahue, L. Anne Hendricks, S. Guadarrama, M. Rohrbach, S. Venugopalan, K. Saenko, and T. Darrell. Long-term recurrent convolutional networks for visual recognition and description. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2625–2634, 2015. 1, 2, 5

[8] S. El Hihi and Y. Bengio. Hierarchical recurrent neural networks for long-term dependencies. In *NIPS*, volume 400, page 409. Citeseer, 1995. 2

[9] A. Erişir, S. C. Van Horn, and S. M. Sherman. Relative numbers of cortical and brainstem inputs to the lateral geniculate nucleus. *Proceedings of the National Academy of Sciences*, 94(4):1517–1520, 1997. 1

[10] D. J. Felleman and D. C. Van Essen. Distributed hierarchical processing in the primate cerebral cortex. *Cerebral cortex*, 1(1):1–47, 1991. 1, 3

[11] D. Fitzpatrick, W. M. Usrey, B. R. Schofield, and G. Einstein. The sublaminar organization of corticogeniculate neurons in layer 6 of macaque striate cortex. *Visual neuroscience*, 11(02):307–315, 1994. 1

[12] F. A. Gers and J. Schmidhuber. Recurrent nets that time and count. In *Neural Networks, 2000. IJCNN 2000, Proceedings of the IEEE-INNS-ENNS International Joint Conference on*, volume 3, pages 189–194. IEEE, 2000. 2

[13] R. Guillery. A quantitative study of synaptic interconnections in the dorsal lateral geniculate nucleus of the cat. *Z. Zellforsch*, 96:39–48, 1969. 1

[14] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *arXiv preprint arXiv:1512.03385*, 2015. 3

[15] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997. 1, 2, 5, 6

[16] Y. Jiang, J. Liu, A. R. Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. Thumos challenge: Action recognition with a large number of classes. In *ECCV Workshop*, 2014. 5

[17] E. R. Kandel, J. H. Schwartz, T. M. Jessell, S. A. Siegelbaum, and A. Hudspeth. *Principles of neural science*, volume 4. McGraw-hill New York, 2000. 2

[18] A. Karpathy, G. Toderici, S. Shetty, T. Leung, R. Sukthankar, and L. Fei-Fei. Large-scale video classification with convolutional neural networks. In *Proceedings of IEEE conference on Computer Vision and Pattern Recognition*, pages 1725–1732, 2014. 8

[19] J. Koutnik, K. Greff, F. Gomez, and J. Schmidhuber. A clockwork rnn. *arXiv preprint arXiv:1402.3511*, 2014. 2

[20] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. Hmdb: a large video database for human motion recognition. In *Proceedings of IEEE International Conference on Computer Vision*, pages 2556–2563, 2011. 5, 8

[21] V. A. Lamme, H. Super, and H. Spekreijse. Feedforward, horizontal, and feedback processing in the visual cortex. *Current opinion in neurobiology*, 8(4):529–535, 1998. 3

[22] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 3

[23] L. Liu, L. Shao, and P. Rockett. Boosted key-frame selection and correlated pyramidal motion-feature representation for human action recognition. *Pattern recognition*, 46(7):1810–1818, 2013. 8

[24] S. L. Macknik and S. Martinez-Conde. The role of feedback in visual attention and awareness. *Cognitive Neurosciences*, 2009. 1

[25] T. Mikolov, A. Joulin, S. Chopra, M. Mathieu, and M. Ranzato. Learning longer memory in recurrent neural networks. *arXiv preprint arXiv:1412.7753*, 2014. 2

[26] Y. Pan, T. Mei, T. Yao, H. Li, and Y. Rui. Jointly modeling embedding and translation to bridge video and language. *arXiv preprint arXiv:1505.01861*, 2015. 1

[27] X. Peng, L. Wang, X. Wang, and Y. Qiao. Bag of visual words and fusion methods for action recognition: Comprehensive study and good practice. *Computer Vision and Image Understanding*, 2016. 8

[28] S. Sharma, R. Kiros, and R. Salakhutdinov. Action recognition using visual attention. *arXiv preprint arXiv:1511.04119*, 2015. 8

[29] S. M. Sherman and R. Guillery. The role of the thalamus in the flow of information to the cortex. *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, 357(1428):1695–1708, 2002. 1

[30] Y. Shi, T. Yonghong, Y. Wang, and T. Huang. Sequential deep trajectory descriptor for action recognition with three-stream cnn. *arXiv preprint arXiv:1609.03056*, 2016. 1, 2, 5

[31] Y. Shi, W. Zeng, T. Huang, and Y. Wang. Learning deep trajectory descriptor for action recognition in videos using deep neural networks. In *Proceedings of IEEE International Conference on Multimedia and Expo*, pages 1–6, 2015. 2

[32] H. T. Siegelmann and E. D. Sontag. On the computational power of neural nets. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 440–449. ACM, 1992. 1

[33] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. In *Advances in Neural Information Processing Systems*, pages 568–576, 2014. 2, 5, 6

[34] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 3

[35] K. Soomro, A. R. Zamir, and M. Shah. Ucf101: A dataset of 101 human actions classes from videos in the wild. *arXiv preprint arXiv:1212.0402*, 2012. 5

[36] L. Sun, K. Jia, D.-Y. Yeung, and B. E. Shi. Human action recognition using factorized spatio-temporal convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4597–4605, 2015. 8

[37] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–9, 2015. 3, 5, 6

[38] G. Varol, I. Laptev, and C. Schmid. Long-term temporal convolutions for action recognition. *arXiv preprint arXiv:1604.04494*, 2016. 8

[39] S. Venugopalan, M. Rohrbach, J. Donahue, R. Mooney, T. Darrell, and K. Saenko. Sequence to sequence-video to text. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4534–4542, 2015. 1

[40] S. Venugopalan, H. Xu, J. Donahue, M. Rohrbach, R. Mooney, and K. Saenko. Translating videos to natural language using deep recurrent neural networks. *arXiv preprint arXiv:1412.4729*, 2014. 1

[41] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 3169–3176, 2011. 2

[42] H. Wang, A. Kläser, C. Schmid, and C.-L. Liu. Dense trajectories and motion boundary descriptors for action recognition. *International Journal of Computer Vision*, 103(1):60–79, 2013. 8

[43] H. Wang and C. Schmid. Action recognition with improved trajectories. In *Proceedings of IEEE International Conference on Computer Vision*, pages 3551–3558, 2013. 2, 8

[44] H. Wang and C. Schmid. Lear-inria submission for the thumos workshop. In *Proceedings of ICCV workshop on action recognition with a large number of classes*, volume 2, page 8, 2013. 8

[45] L. Wang, Y. Qiao, and X. Tang. Motionlets: Mid-level 3d parts for human motion recognition. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 2674–2681, 2013. 8

[46] L. Wang, Y. Qiao, and X. Tang. Action recognition with trajectory-pooled deep-convolutional descriptors. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 4305–4314, 2015. 2, 8

[47] L. Wang, Y. Qiao, and X. Tang. Mofap: A multi-level representation for action recognition. *International Journal of Computer Vision*, pages 1–18, 2015. 8

[48] L. Wang, Y. Xiong, Z. Wang, and Y. Qiao. Towards good practices for very deep two-stream convnets. *arXiv preprint arXiv:1507.02159*, 2015. 2, 5, 7, 8

[49] L. Wang, Y. Xiong, Z. Wang, Y. Qiao, D. Lin, X. Tang, and L. Van Gool. Temporal segment networks: Towards good practices for deep action recognition. *arXiv preprint arXiv:1608.00859*, 2016. 5, 8

[50] P. J. Werbos. Backpropagation through time: what it does and how to do it. *Proceedings of the IEEE*, 78(10):1550–1560, 1990. 3

[51] Z. Wu, X. Wang, Y.-G. Jiang, H. Ye, and X. Xue. Modeling spatial-temporal clues in a hybrid deep learning framework for video classification. In *Proceedings of the 23rd ACM international conference on Multimedia*, pages 461–470, 2015. 2, 7, 8

[52] L. Yao, A. Torabi, K. Cho, N. Ballas, C. Pal, H. Larochelle, and A. Courville. Describing videos by exploiting temporal structure. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4507–4515, 2015. 1

[53] J. Yue-Hei Ng, M. Hausknecht, S. Vijayanarasimhan, O. Vinyals, R. Monga, and G. Toderici. Beyond short snippets: Deep networks for video classification. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition*, pages 4694–4702, 2015. 1, 2, 8

[54] C. Zach, T. Pock, and H. Bischof. A duality based approach for realtime tv-l 1 optical flow. In *Pattern Recognition*, pages 214–223. Springer, 2007. 5